



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
09/496,844	02/02/2000	Patrick Knebel	10971393-1	6757

22879 7590 04/16/2004

HEWLETT PACKARD COMPANY
P O BOX 272400, 3404 E. HARMONY ROAD
INTELLECTUAL PROPERTY ADMINISTRATION
FORT COLLINS, CO 80527-2400

EXAMINER

HUISMAN, DAVID J

ART UNIT	PAPER NUMBER
----------	--------------

2183

DATE MAILED: 04/16/2004

10

Please find below and/or attached an Office communication concerning this application or proceeding.

24

Office Action Summary

Application No.

09/496,844

Applicant(s)

KNEBEL ET AL.

Examiner

David J. Huisman

Art Unit

2183

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If the period for reply specified above is less than thirty (30) days, a reply within the statutory minimum of thirty (30) days will be considered timely.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) ☒ Responsive to communication(s) filed on 24 March 2004.
- 2a) ☒ This action is **FINAL**. 2b) ☐ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) ☒ Claim(s) 1,3-5,7-15,18,19,21,22 and 25-29 is/are pending in the application.
- 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) ☐ Claim(s) _____ is/are allowed.
- 6) ☒ Claim(s) 1,3-5,7-15,18,19,21,22 and 25-29 is/are rejected.
- 7) ☐ Claim(s) _____ is/are objected to.
- 8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) ☐ The specification is objected to by the Examiner.
- 10) ☒ The drawing(s) filed on 24 March 2004 is/are: a) ☐ accepted or b) ☒ objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) ☐ All b) ☐ Some * c) ☐ None of:
1. ☐ Certified copies of the priority documents have been received.
2. ☐ Certified copies of the priority documents have been received in Application No. _____.
3. ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- 1) ☒ Notice of References Cited (PTO-892)
- 2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
- 3) ☐ Information Disclosure Statement(s) (PTO-1449 or PTO/SB/08)
Paper No(s)/Mail Date _____
- 4) ☐ Interview Summary (PTO-413)
Paper No(s)/Mail Date. _____
- 5) ☐ Notice of Informal Patent Application (PTO-152)
- 6) ☐ Other: _____

Art Unit: 2183

DETAILED ACTION

1. Claims 1, 3-5, 7-15, 18-19, 21-22, and 25-29 have been examined.

Papers Submitted

2. It is hereby acknowledged that the following papers have been received and placed of record in the file: #18. Amendment "E" as received on 3/24/2004.

Drawings

3. The drawings are objected to because of the following minor informalities: Regarding Fig.5, the figure should be labeled as Fig.5. In addition, in Fig.5, the text in the boxes should be centered such that the text does not overlap with the box boundaries. A proposed drawing correction or corrected drawings are required in reply to the Office action to avoid abandonment of the application. The objection to the drawings will not be held in abeyance.

Response to Amendment

4. The declaration filed on March 24, 2004 under 37 CFR 1.131 has been considered but is ineffective to overcome the Col et al. reference (6,330,657).
5. The evidence submitted is insufficient to establish a reduction to practice of the invention in this country or a NAFTA or WTO member country prior to the effective date of the Col reference. According to MPEP §715.07, "In general, proof of actual reduction to practice requires a showing that the apparatus actually existed and worked for its intended purpose. However, "there are some devices so simple that a mere construction of them is all that is

Art Unit: 2183

necessary to constitute reduction to practice.” In re Asahi /America Inc., 94-1249 (Fed. Cir. 1995) (Citing Newkirk v. Lulegian, 825 F.2d 1581, 3USPQ2d 1793 (Fed. Cir. 1987) and Sachs v. Wadsworth, 48 F.2d 928, 929, 9 USPQ 252, 253 (CCPA 1931).” The examiner concedes that the RTL source code (Exhibits 1 and 2) may be enough to show that the apparatus actually existed before the Col reference, but the exhibits are not simple enough to show that the apparatus operated as intended. Therefore, in order to establish an actual reduction to practice, the applicant must submit testing evidence and simulations of the RTL code showing that the apparatus did in fact operate as intended. Since no such showing has been provided, applicant has failed to establish actual reduction to practice. See MPEP §2138.05 for more details.

6. Since applicant has failed to establish an actual reduction of practice, applicant is required to show conception of the invention prior to the effective date of the reference coupled with due diligence from prior to the reference date to the filing date of the application (constructive reduction to practice). See MPEP §715.07. However, the evidence submitted is insufficient to establish diligence from a date prior to the date of reduction to practice of the Col reference to a constructive reduction to practice. Consequently, the declaration file under 37 CFR 1.131 is improper, thereby failing to overcome the Col reference.

Claim Objections

7. Claim 1 is objected to because of the following informalities: The examiner believes the word --the-- should be inserted before “at” in line 5. Appropriate correction is required.

8. Claim 12 is objected to because of the following informalities: The examiner believes the word --the-- should be inserted before “at” in line 11. Appropriate correction is required.

Art Unit: 2183

Withdrawn Rejections

9. Through amendment, applicant has overcome the rejections set forth in the Office Action mailed on December 24, 2003, for claims 1, 3-5, 7-15, 18, 19, and 21-24. However, upon further consideration, a new ground(s) of rejection has been made below.

Claim Rejections - 35 USC § 102

10. The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(f) he did not himself invent the subject matter sought to be patented.

11. Claims 1, 3-5, 7-15, 18-19, 21-22, and 25-29 are rejected under 35 U.S.C. 102(f) because the applicant did not invent the claimed subject matter. The instant application (09/496,844) has been filed by inventors Patrick Knebel and Kevin Safford. However, the submission of the RTL code listings on March 24, 2004, seems to show otherwise. Firstly, while Patrick Knebel's name does appear in the RTL code listings, the examiner has been unable to find Kevin Safford's name anywhere in the RTL code listings. Secondly, Exhibit 2 (RTL code listing dated December 4, 2008) includes a code revision history section (pages 13-16) which shows that the following people have modified/revised the code: Rohit Bhatia (appears 18 times), Gary Welte (appears 5 times), Michael J. Lee (appears 2 times), Generic CAD User (appears 13 times), Patrick Knebel (appears 1 time), Ravi Koshy (appears 3 times), and Preston Renstrom (appears 1 time). Thus, it is not clear: a) how Kevin Safford is involved, and b) if these additional people are also inventors of this invention since they had modified the code at one or more points in time.

Claim Rejections - 35 USC § 103

12. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

13. Claims 1, 3, 4, 7, 8, 12-15, and 19 are rejected under 35 U.S.C. 103(a) as being unpatentable over Col et al., U.S. Patent No. 6,330,657 B1 (as applied in the previous Office Action and herein referred to as Col) in view of Nakajima, U.S. Patent No. 5,537,561, in view of Shang et al., U.S. Patent No. 5,764,971 (as applied in the previous Office Action and herein referred to as Shang), and further in view of Song, U.S. Patent No. 5,546,599.

14. Referring to claim 1, Col has taught a method for processing software instructions comprising:

a) decomposing a macroinstruction into a plurality of microinstructions. See Fig.4, steps 402 and 404.

b) Col has not taught determining whether at least two of the plurality of microinstructions are required to issue in parallel. However, Nakajima has taught such a concept. See Fig.12(a) and Fig.12(b). It should first be realized that when the first stages of both pipelines (81 and 91) are free and bypassing will not occur, the two instructions held in instruction memory 60 are required to issue in parallel. That is, they are to issue in parallel, as opposed to the scenario shown in Fig.11(a), Fig.11(b), Fig.13(a), and Fig.13(b) where the instructions in instruction memory are not required to issue in parallel due to resource conflicts. For the instructions in

Art Unit: 2183

Fig. 12(a) and Fig. 12(b), it is beneficial to require them to issue in parallel because they are both writing to the same register. Instead of spending more time executing them sequentially, Nakajima requires them to issue in parallel and allows the second one to write the result, thereby speeding up execution. Overall, it would have been obvious to one of ordinary skill in the art at the time of the invention to modify Col's system so that it determines whether the microinstructions are required to issue in parallel because the system would gain the advantage of increased execution speed when the instructions can be issued in parallel, while still allowing instructions to not be issued in parallel when resource conflicts occur.

c) Col has taught forcing the parallel issue of at least two of the plurality of microinstructions simultaneously. See column 3, lines 52-56. Col has not taught forcing the parallel issue regardless of conflict checking. However, Nakajima has taught such a concept. Looking at Fig. 12(a), it should be noted that both registers write to the same destination (f00), i.e. a WAW hazard. Nakajima's system merely detects this conflict and adds tag information to each instruction, as shown in Fig. 12(b). These tags are used to correct the conflict at a later point. However, as seen from Fig. 12(b), the conflicting instructions are still issued in parallel. A description of this process is given in columns 11 and 12. The reason this is done is to realize parallel processing with a small amount of hardware without deteriorating processing efficiency though data dependencies are secured. See column 13, lines 33-38. This hardware reduction comes from the lack of need for waiting buffers as described in column 2, lines 46-51. And one of ordinary skill in the art would have realized that the more operations that are issued (and executed in parallel, the faster the system). Therefore, in order to realize parallel processing (which results in higher throughput) while reducing hardware and recognizing dependencies, it

Art Unit: 2183

would have been obvious to one of ordinary skill in the art at the time of the invention to force a parallel issue regardless of conflict checking.

d) executing the at least two microinstructions simultaneously, in lockstep using functional units in a floating-point unit. See column 3, lines 31-35, and Fig.6 (note in cycle 7 that two microinstructions are executed in parallel). Furthermore, note from column 20, lines 32-38 and note that this parallel execution can occur using multiple floating-point functional units.

e) Col has not explicitly taught:

e1) determining whether an exception occurs in any of the microinstructions, before writing results of the executing to result registers, wherein the determining step is performed prior to any writing step.

e2) if an exception occurs in any of the microinstructions, canceling all of the microinstructions and preventing the results of the executing from being written to the result registers.

e3) if no exception occurs in any of the microinstructions, writing the results of the executing to the result registers.

However, exceptions and the advantages of detecting exceptions are well known, accepted, and expected to occur in the art. In general, an exception is an interruption to the normal flow of program control, caused by the program itself or by executing an illegal instruction. Shang has taught a system in which macroinstructions are translated into a plurality of microinstructions and if an exception is detected within any one of those microinstructions, then the rest of the microinstructions are cancelled and results are not written to the result registers. Otherwise, if no exception is detected, then the results of the microinstructions are written to the result registers.

Art Unit: 2183

See Fig.6 and note that if exceptions (interrupts) have been detected in at least one of the microinstructions at step 104, then the system is flushed (cancellation of each microinstruction) at step 108. If no exceptions were detected at step 104, then the results are written to the result registers at step 106. From this flowchart it can be seen that the determining of an exception occurs before the final writing step (step 106). The final writing step is "any writing step" as claimed by applicant. More specifically, the examiner is reading the claim as saying "the determining step is performed prior to a final writing step" because a final writing step is any writing step. In Col's system, since a macroinstruction is broken up into microinstructions, an exception in a single microinstruction would mean that an exception has occurred in the overall macroinstruction, and therefore, all of the microinstructions that represent a single macroinstruction, should not be able to change the state of the system. An advantage of Shang's scheme would be to avoid having to undo the changes made by the undesired execution of a microinstruction that is part of a faulty macroinstruction. This will prevent a reduction in throughput in that the extra time required to perform an undo-operation would not be necessary. Therefore, in order to maximize the efficiency of the overall system, it would have been obvious to one of ordinary skill in the art at the time of the invention to modify Col's system such that it employs exception detection among microinstructions as taught by Shang.

e4) Shang has not taught that the method does not write any results to temporary registers. However, Song has taught detecting exceptions by instructions without any writing occurring to temporary storage. See column 15, lines 37-51. More specifically, instruction results are written directly to final storage so that precise interrupts and precise exceptions will be achieved. In addition, not writing to temporary storage results

Art Unit: 2183

in the removal of an extra writing step, thereby allowing for faster execution. As a result, it would have been obvious to one of ordinary skill in the art at the time of the invention to modify Shang such that temporary storage is not written to when detecting an exception.

15. Referring to claim 3, Col in view of Nakajima in view of Shang and further in view of Song has taught a method as described in claim 1. Col has further taught that the microinstructions are executed on separate execution units, but appear as though they were executed on a single execution unit. See column 20, lines 32-38. Note that multiple floating-point execution units are used per clock cycle in order to execute multiple microinstructions. For example, the microinstructions of cycle 7 in Fig.6 must be executed on different execution units if they are executed in parallel. Then the results of each microinstruction are written to the appropriate storage via store logic (component 420 of Fig.4). See column 17, lines 3-16. Note that the store logic retrieves all of the results from each microinstruction execution and writes the data to the appropriate place. Finally, from Fig.6 (cycle 7), it should be realized that the separate, but parallel, execution of LD T1,[BX] and ADD AX,T1, will produce a result that is expected for the ADD AX,[BX] macroinstruction.

16. Referring to claim 4, Col in view of Nakajima in view of Shang and further in view of Song has taught a method as described in claim 1. Col has further taught that all of the microinstructions are executed on the same clock cycle. See Fig.6, cycle 7, for instance.

17. Referring to claim 7, Col in view of Nakajima in view of Shang and further in view of Song has taught a method as described in claim 1. Furthermore, Col has taught that the system

Art Unit: 2183

allows a single instruction to operate on multiple single-precision floating-point values. See column 3, lines 61-63.

18. Referring to claim 8, Col in view of Nakajima in view of Shang and further in view of Song has taught a method as described in claim 1. Col has not explicitly taught that a flag is updated based upon a result of the execution of the microinstructions. However, it is well known, accepted, and expected in the art that processors contain a status register. The status register contains bits that are set or cleared based on the result of an operation. Some of the more common flags are ones that indicate a result of zero, a negative number, and overflow. These flags can then be checked in conditional situations, such as branches. Therefore, it would have been obvious to one of ordinary skill in the art to update a flag based upon the result of the execution of the microinstructions.

19. Referring to claim 12, Col has taught a computer system comprising:

a) a processor comprising:

a1) a floating-point unit comprising a plurality of functional units adapted to execute microinstructions. See Fig.4, component 414, and column 20, lines 32-38.

a2) Col has not explicitly taught a computer system with a ROM. However, Official Notice is taken that ROMs are well known, accepted, and expected in the art. Processors contain Read-Only Memory to store essential software of the computer. Because it's non-volatile memory, ROM does not lose its contents when the power is turned off. Therefore, a ROM chip is used to store control programs for the computer, such as the bootstrap program (which tells the computer how to start and load the operating system) and other types of configuration information. As a result, it would have been obvious to

Art Unit: 2183

one of ordinary skill in the art at the time of the invention to provide some type of Rom in Col's system.

a3) a plurality of floating-point registers. See column 5, lines 49-52.

b) wherein the processor is configured to emulate an instruction set by:

b1) decomposing a macroinstruction into a plurality of microinstructions. See Fig. 4, steps 402 and 404.

b2) Col has not taught determining whether at least two of the plurality of microinstructions are required to issue in parallel. However, Nakajima has taught such a concept. See Fig. 12(a) and Fig. 12(b). It should first be realized that when the first stages of both pipelines (81 and 91) are free and bypassing will not occur, the two instructions held in instruction memory 60 are required to issue in parallel. That is, they are to issue in parallel, as opposed to the scenario shown in Fig. 11(a), Fig. 11(b), Fig. 13(a), and Fig. 13(b) where the instructions in instruction memory are not required to issue in parallel due to resource conflicts. For the instructions in Fig. 12(a) and Fig. 12(b), it is beneficial to require them to issue in parallel because they are both writing to the same register. Instead of spending more time executing them sequentially, Nakajima requires them to issue in parallel and allows the second one to write the result, thereby speeding up execution. Overall, it would have been obvious to one of ordinary skill in the art at the time of the invention to modify Col's system so that it determines whether the microinstructions are required to issue in parallel because the system would gain the advantage of increased execution speed when the instructions can be issued in parallel, while still allowing instructions to not be issued in parallel when resource conflicts occur.

Art Unit: 2183

b3) forcing the parallel issue of at least two of the plurality microinstructions simultaneously to the functional units. See column 3, lines 52-56. Col has not taught forcing the parallel issue regardless of conflict checking. However, Nakajima has taught such a concept. See Fig. 12(a) and Fig. 12(b). Note in Fig. 12(a) that the two instructions in storage 60 are to be issued in parallel to pipelines 80 and 90. It should be noted that both registers write to the same destination (froo), i.e. a WAW hazard. Nakajima's system merely detects this conflict and adds tag information to each instruction, as shown in Fig. 12(b). These tags are used to correct the conflict at a later point. However, as seen from Fig. 12(b), the conflicting instructions are still issued in parallel. A description of this process is given in columns 11 and 12. The reason this is done is to realize parallel processing with a small amount of hardware without deteriorating processing efficiency though data dependencies are secured. See column 13, lines 33-38. This hardware reduction comes from the lack of need for waiting buffers as described in column 2, lines 46-51. And one of ordinary skill in the art would have realized that the more operations that are issued (and executed in parallel, the faster the system). Therefore, in order to realize parallel processing (which results in higher throughput) while reducing hardware and recognizing dependencies, it would have been obvious to one of ordinary skill in the art at the time of the invention to force a parallel issue regardless of conflict checking.

b4) Col has not explicitly taught determining whether an exception occurs in any of the functional units wherein the determining step is performed prior to any setting step and the method does not set any temporary registers, setting result registers for results of each of the functional units only if no exception occurs in any of the functional units, and if an

exception occurs in any of the microinstructions, canceling all of the microinstructions and preventing the setting of result registers for all of the functional units. However, exceptions and the advantages of detecting exceptions are well known, accepted, and expected to occur in the art. In general, an exception is an interruption to the normal flow of program control, caused by the program itself or by executing an illegal instruction. Shang has taught a system in which macroinstructions are translated into a plurality of microinstructions and if an exception is detected within any one of those microinstructions, then the rest of the microinstructions are cancelled and results are not written to the result registers. Otherwise, if no exception is detected, then the results of the microinstructions are written to the result registers. See Fig.6 and note that if exceptions (interrupts) have been detected in at least one of the microinstructions at step 104, then the system is flushed (cancellation of each microinstruction) at step 108. If no exceptions were detected at step 104, then the results are written to the result registers at step 106. From this flowchart it can be seen that the determining of an exception occurs before the final writing (setting) step (step 106). The final setting step is "any setting step" as claimed by applicant. More specifically, the examiner is reading the claim as saying "the determining step is performed prior to a final setting step" because a final setting step is any setting step. In Col's system, since a macroinstruction is broken up into microinstructions, an exception in a single microinstruction would mean that an exception has occurred in the overall macroinstruction, and therefore, all of the microinstructions that represent a single macroinstruction, should not be able to change the state of the system. An advantage of Shang's scheme would be to avoid having to

undo the changes made by the undesired execution of a microinstruction that is part of a faulty macroinstruction. This will prevent a reduction in throughput in that the extra time required to perform an undo-operation would not be necessary. Therefore, in order to maximize the efficiency of the overall system, it would have been obvious to one of ordinary skill in the art at the time of the invention to modify Col's system such that it employs exception detection among microinstructions as taught by Shang.

b5) Shang has not taught that the method does not set any temporary registers. However, Song has taught detecting exceptions by instructions without setting temporary storage. See column 15, lines 37-51. More specifically, instruction results are written directly to final storage so that precise interrupts and precise exceptions will be achieved. In addition, not setting temporary storage results in the removal of an extra step (the setting step), thereby allowing for faster execution. As a result, it would have been obvious to one of ordinary skill in the art at the time of the invention to modify Shang such that temporary storage is not set when detecting an exception.

20. Referring to claim 13, Col in view of Nakajima in view of Shang and further in view of Song has taught a computer system as described in claim 12. Col has further taught that the processor is further configured to emulate the instruction set by executing all of the microinstructions. See column 3, lines 31-35, and Fig.6 (note in cycle 7 that two microinstructions are executed in parallel).

21. Referring to claim 14, Col in view of Nakajima in view of Shang and further in view of Song has taught a computer system as described in claim 13. Furthermore, it has been noted that

Art Unit: 2183

the computer system of claim 14 performs the method of claim 3. Therefore, claim 14 is rejected for the same reasons set forth in the rejection of claim 3 above.

22. Referring to claim 15, Col in view of Nakajima in view of Shang and further in view of Song has taught a computer system as described in claim 14. Furthermore, it has been noted that the computer system of claim 15 performs the method of claim 8. Therefore, claim 15 is rejected for the same reasons set forth in the rejection of claim 8 above.

23. Referring to claim 19, Col in view of Nakajima in view of Shang and further in view of Song has taught a method as described in claim 1.

a) Col has further taught that the step of issuing comprises forcing the microinstructions to issue simultaneously, in lockstep with each other. See column 3, lines 52-56.

b) Col has not explicitly taught that the step of canceling comprises canceling all of the plurality of microinstructions without regard to the relative ages of the microinstructions and without using a backoff mechanism. However, recall from the rejection of claim 1 above, that it would have been an obvious modification to add exception detection functionality to Col's system (in view of Shang). This modification, as described above, would allow for the detection of an exception within a single microinstruction and if an exception is triggered, the other related microinstructions will be subject to cancellation. Furthermore, recall that if interrupts are detected within Shang's system, results are not written to result registers, and consequently, there is no need for a backoff mechanism to undo results that have been incorrectly written.

24. Claim 5 is rejected under 35 U.S.C. 103(a) as being unpatentable over Col in view of Nakajima in view of Shang in view of Song, as applied above, and further in view of Hennessy

Art Unit: 2183

and Patterson, Computer Architecture - A Quantitative Approach, 2nd Edition, 1996 (as applied in the previous Office Action and herein referred to as Hennessy).

25. Referring to claim 5, Col in view of Nakajima in view of Shang in view of Song has taught a method as described in claim 1. Col has further taught that his system can execute floating-point operations, which is supported in column 20, lines 32-38. Col has not explicitly taught that the microinstructions are executed over multiple clock cycles. However, it is well known, accepted, and expected in the art that floating-point operations can consume more than one clock cycle for execution purposes. Hennessy has disclosed this concept on page 187. Hennessy has also shown a pipeline that accommodates floating-point execution through multiple execution stages. See Fig.3.44 and Fig.3.45 on page 190. Since it has been disclosed by Hennessy that a floating-point operation takes multiple instruction execution cycles, it follows that it would have been obvious to one of ordinary skill in the art at the time of the invention to use a pipelined floating-point execution unit with multiple execution stages if floating-point operations are desired.

26. Claims 10, 11, 21, and 22 are rejected under 35 U.S.C. 103(a) as being unpatentable over Col in view of Nakajima in view of Shang in view of Song, as applied above, and further in view of Intel ®, Intel ® Architecture Optimization Reference Manual, 1998-1999 (as applied in the previous Office Action and herein referred to as Intel).

27. Referring to claim 10, Col has taught a method for processing software instructions comprising:

Art Unit: 2183

a) providing two microinstructions to emulate a high-half and a low-half Single Instruction Multiple-Data Extensions (SSE) operation. See Fig.6 and note in cycle 7 that two microinstructions are executed in parallel. Col has not explicitly taught that the operation is an SSE operation. However, Intel has taught that SSE instructions are used to accelerate performance of applications regarding 3D geometry. See page 1-12. As a result, in order to increase system performance, it would have been obvious to one of ordinary skill in the art at the time of the invention to allow the system of Col to emulate SSE operations.

b) Col has not taught determining whether the two microinstructions are required to issue in parallel. However, Nakajima has taught such a concept. See Fig.12(a) and Fig.12(b). It should first be realized that when the first stages of both pipelines (81 and 91) are free and bypassing will not occur, the two instructions held in instruction memory 60 are required to issue in parallel. That is, they are to issue in parallel, as opposed to the scenario shown in Fig.11(a), Fig.11(b), Fig.13(a), and Fig.13(b) where the instructions in instruction memory are not required to issue in parallel due to resource conflicts. For the instructions in Fig.12(a) and Fig.12(b), it is beneficial to require them to issue in parallel because they are both writing to the same register. Instead of spending more time executing them sequentially, Nakajima requires them to issue in parallel and allows the second one to write the result, thereby speeding up execution. Overall, it would have been obvious to one of ordinary skill in the art at the time of the invention to modify Col's system so that it determines whether the microinstructions are required to issue in parallel because the system would gain the advantage of increased execution speed when the instructions can be issued in parallel, while still allowing instructions to not be issued in parallel when resource conflicts occur.

Art Unit: 2183

c) forcing the high-half and low-half operations to issue in parallel. See column 3, lines 52-56.

Col has not taught forcing the parallel issue regardless of conflict checking. However, Nakajima has taught such a concept. See Fig. 12(a) and Fig. 12(b). Note in Fig. 12(a) that the two instructions in storage 60 are to be issued in parallel to pipelines 80 and 90. It should be noted that both registers write to the same destination (f00), i.e. a WAW hazard. Nakajima's system merely detects this conflict and adds tag information to each instruction, as shown in Fig. 12(b). These tags are used to correct the conflict at a later point. However, as seen from Fig. 12(b), the conflicting instructions are still issued in parallel. A description of this process is given in columns 11 and 12. The reason this is done is to realize parallel processing with a small amount of hardware without deteriorating processing efficiency though data dependencies are secured. See column 13, lines 33-38. This hardware reduction comes from the lack of need for waiting buffers as described in column 2, lines 46-51. And one of ordinary skill in the art would have realized that the more operations that are issued (and executed in parallel, the faster the system). Therefore, in order to realize parallel processing (which results in higher throughput) while reducing hardware and recognizing dependencies, it would have been obvious to one of ordinary skill in the art at the time of the invention to force a parallel issue regardless of conflict checking.

d) dispatching the high-half and low-half operations simultaneously to a first floating point unit and to a second floating point unit, respectively. See column 3, lines 52-56, and column 20, lines 32-38.

e) executing the high-half and low-half operations simultaneously, in lockstep. See column 3, lines 31-35, and Fig. 6 (note in cycle 7 that two microinstructions are executed in parallel).

Art Unit: 2183

f) generating a signal from an emulator's hardware. Signals are inherently generated within a computer system. For instance, a clock signal is a basic signal that synchronizes the many different information-processing tasks assigned to the chip. Also, as instructions are fetched and decoded, signals are sent to the appropriate functional units in order to "specify" which operations are to be performed based on the type of instruction.

g) sending the signal to the first and second floating point functional units. Again, the appropriate signals would have to be supplied to the appropriate functional units in order to perform the desired operation.

h) Col has not explicitly taught determining whether an exception is taken in either the first or the second floating point unit, wherein the determining step is performed prior to any writing step and if the exception is taken in either the first or second floating point unit, preventing results from the high-half and low-half operations from being written to result registers, and canceling both the high-half and low-half operations. However, exceptions and the advantages of detecting exceptions are well known, accepted, and expected to occur in the art. In general, an exception is an interruption to the normal flow of program control, caused by the program itself or by executing an illegal instruction. Shang has taught a system in which macroinstructions are translated into a plurality of microinstructions and if an exception is detected within any one of those microinstructions, then the rest of the microinstructions are cancelled and results are not written to the result registers. Otherwise, if no exception is detected, then the results of the microinstructions are written to the result registers. See Fig.6 and note that if exceptions (interrupts) have been detected in at least one of the microinstructions at step 104, then the system is flushed (cancellation of each microinstruction) at step 108. If no exceptions were

Art Unit: 2183

detected at step 104, then the results are written to the result registers at step 106. From this flowchart it can be seen that the determining of an exception occurs before the final writing step (step 106). The final writing step is "any writing step" as claimed by applicant. More specifically, the examiner is reading the claim as saying "the determining step is performed prior to a final writing step" because a final writing step is any writing step. In Col's system, since a macroinstruction is broken up into microinstructions, an exception in a single microinstruction would mean that an exception has occurred in the overall macroinstruction, and therefore, all of the microinstructions that represent a single macroinstruction, should not be able to change the state of the system. An advantage of Shang's scheme would be to avoid having to undo the changes made by the undesired execution of a microinstruction that is part of a faulty macroinstruction. This will prevent a reduction in throughput in that the extra time required to perform an undo-operation would not be necessary. Therefore, in order to maximize the efficiency of the overall system, it would have been obvious to one of ordinary skill in the art at the time of the invention to modify Col's system such that it employs exception detection among microinstructions as taught by Shang.

i) Shang has not taught that the method does not write any results to temporary registers.

However, Song has taught detecting exceptions by instructions without any writing occurring to temporary storage. See column 15, lines 37-51. More specifically, instruction results are written directly to final storage so that precise interrupts and precise exceptions will be achieved. In addition, not writing to temporary storage results in the removal of an extra writing step, thereby allowing for faster execution. As a result, it would have been obvious to one of ordinary skill in

Art Unit: 2183

the art at the time of the invention to modify Shang such that temporary storage is not written to when detecting an exception.

j) Col has not explicitly taught updating MXCSR flags based upon the results of the first and second floating point units. However, the general idea of a status register is well known, accepted, and expected in the art. The MXCSR register contains flags that are common to other processor status registers. These bits (flags) are set and cleared based on results from operations. For instance, if the result of an addition were zero, a flag indicating a zero-result would be set in the status register. Or, perhaps an overflow occurred. A flag in the status register would be set to specify that as well. Conditional statements such as branches then reference these flags in order to determine the direction of the program. Therefore, in order to provide a readable status of the processor so that programs (via branches) flow according to previously obtained results, it would have been obvious to one of ordinary skill in the art at the time of the invention to update flags that are found in the MXCSR register.

28. Referring to claim 11, Col in view of Nakajima in view of Shang in view of Song and further in view of Intel has taught a method as described in claim 10. Recall that it has been established that it would have been obvious to one of ordinary skill in the art at the time of the invention to modify Col such that it can detect and correct exceptions in a manner taught by Shang. Note also that Shang the preventing and canceling steps are performed regardless of the relative ages of the microinstructions. See Fig.6.

29. Referring to claim 21, Col in view of Nakajima in view of Shang and further in view of Song has taught a method as described in claim 1 wherein the step of executing comprises executing using a plurality of functional units of a floating-point unit. Col in view of Nakajima

Art Unit: 2183

in view of Shang in view of Song has not explicitly taught the emulation of Streaming Single Instruction Multiple-Data Extensions (SSE) instructions. However, as discussed above, Intel has taught that SSE instructions are used to accelerate performance of applications regarding 3D geometry. See page 1-12. As a result, in order to increase system performance, it would have been obvious to one of ordinary skill in the art at the time of the invention to allow for the emulation of SSE instructions within the system of Col. Furthermore, it is inherent that within computer systems, as instructions are fetched and decoded, signals are sent to the appropriate functional units in order to “specify” which operations are to be performed based on the type of instruction. Therefore, if Col’s system included SSE instructions, as established above, Col would have inherently taught:

- a) generating a signal via hardware that indicates that the functional units are emulating an SSE instruction and sending the signal to the functional units. Again, the appropriate signals would have to be supplied to the appropriate functional units in order to perform the desired operation.
- b) Also, it is inherent that determining an exception would occur after the functional unit has received its signal. For instance, the only way an overflow exception could be detected is by checking the result of an operation, which would only be obtained subsequent to “telling” the functional unit which operation to perform.

30. Referring to claim 22, Col in view of Nakajima in view of Shang in view of Song has taught a system as described in claim 12. Furthermore, it has been noted that the system of claim 22 performs the method of claim 21. Therefore, claim 22 is rejected for the same reasons set forth in the rejection of claim 21.

Art Unit: 2183

31. Claim 9 is rejected under 35 U.S.C. 103(a) as being unpatentable over Col in view of Nakajima in view of Shang in view of Song, as applied above, and further in view of Phillips et al., U.S. Patent No. 6,038,652 (as applied in the previous Office Action and herein referred to as Philips).

32. Referring to claim 9, Col in view of Nakajima in view of Shang in view of Song has taught a method as described in claim 1.

a) Col has not explicitly taught that if an unmasked exception occurs, canceling the execution of all of the plurality of microinstructions, without regard to the relative ages of each of the plurality of microinstructions, and invoking a microcode handler. However, recall from the rejection of claim 1 above, that it would have been an obvious modification to add exception detection functionality to Col's system (in view of Shang). This modification, as described above, would allow for the detection of an exception within a single microinstruction and if an exception is triggered, the other related microinstructions would be subject to cancellation. In addition, Shang has taught that the triggering of an exception will result in the invocation of a microcode handler (referred to as an interrupt service routine). See column 11, lines 21-25. In general, the handler is invoked in order to correct the cause and effects of the exception and allow the processor to continue execution. Therefore, in order to correctly service an exception, it would have been obvious to one of ordinary skill in the art at the time of the invention to implement a microcode handler, which must be invoked upon exception detection.

b) Col in view of Shang has not explicitly taught updating at least one exception flag (when an unmasked exception occurs) by independently generating a logical OR of exceptions for a plurality of functional units. However, Phillips has taught the concept of simultaneously

Art Unit: 2183

checking SIMD elements for exceptions and combining each individual exception into an overall exception. See FIG.2. Furthermore, the combining element (230) in FIG.2 can be implemented as an OR gate that generates a flag (240) used to specify whether or not an exception has occurred. See column 3, lines 60-63. A person of ordinary skill in the art would have recognized that the concept of Phillips would be applicable in Col's system in order to check for exceptions during the parallel execution of microinstructions. In a SIMD processor (as taught by Col), the overhead incurred to process the many possible exceptions generated by SIMD elements may be expensive and lead to degradation in performance. The system of Philips provides an efficient technique to report exceptions occurring in computing complex functions on a SIMD machine. An advantage to this scheme is that since an exception flag is produced according to a parallel execution of microinstructions as opposed to a serial execution of microinstructions, the processor will be able to detect an exception sooner and therefore sooner make the determination that the involved microinstructions should be cancelled. Therefore, it would have been obvious to one of ordinary skill in the art at the time of the invention to update at least one exception flag in Col's system based on the exception check for a plurality of microinstructions.

33. Claim 18 is rejected under 35 U.S.C. 103(a) as being unpatentable over Col in view of Nakajima in view of Shang in view of Song and further in view of Intel, as applied above, and further in view of Makineni et al., U.S. Patent No 6,321,327 B1 (herein referred to as Makineni).

34. Referring to claim 18, Col in view of Nakajima in view of Shang in view of Song has taught a computer system as described in claim 12. Col has further taught the general use of

Art Unit: 2183

SIMD instructions, which is the format used by SSE instructions. See column 3, lines 61-63.

However, Col has not disclosed the specific use of Streaming Single Instruction Multiple-Data Extensions (SSE) instructions. Intel has taught that SSE instructions are used to accelerate performance of applications regarding 3D geometry. See page 1-12. As a result, in order to increase system performance, it would have been obvious to one of ordinary skill in the art at the time of the invention to implement an SSE instruction set within the system of Col. In addition, Col has taught that the SIMD execution units perform operations (such as an add) on multiple operands from a first SIMD register with corresponding multiple operands from a second SIMD register. See column 12, lines 1-8. Col has not explicitly taught the use of two 82-bit floating point registers for emulating four 32-bit single-precision floating-point values in an SSE register. However, Makineni has taught the use of 82-bit registers to hold 32-bit single precision floating-point numbers. See Fig.2B and Fig.3. Since SIMD operations involve performing a single operation on multiple pairs of data elements, it follows that multiple pairs of operands must be available to the SIMD execution unit. A person of ordinary skill in the art would have recognized that by implementing 82-bit registers, multiple pairs of operands would be supplied to a SIMD execution unit by using multiple registers. In addition, by packing more than one data element into a single register, the amount of addressable registers could be decreased, resulting in less wires used for addressing purposes. Finally, each of the standard IEEE floating-point formats can be specified through the use of a single 82-bit floating point register, allowing the system to operate on different-precision operands depending on the situation. See the floating-point standards on page A-13 of Hennessy and note Fig.2A of Makineni shows a double-extended precision floating-point number. Therefore, in order to decrease the amount of

Art Unit: 2183

hardware, while assuring the system has the capability of processing a wide variety of floating-point numbers, it would have been obvious to one of ordinary skill in the art to use two 82-bit floating point registers for emulating four 32-bit single-precision floating-point values in an SSE register.

35. Claim 25 is rejected under 35 U.S.C. 103(a) as being unpatentable over Col in view of Nakajima in view of Shang in view of Song, as applied above, and further in view of Blomgren, U.S. Patent No. 5,598,546.

36. Referring to claim 25, Col in view of Nakajima in view of Shang in view of Song has taught a method as described in claim 1. Col in view of Nakajima in view of Shang in view of Song has not taught delaying the issue of a first microinstruction if it is determined that said first microinstruction is required to be issued in parallel with at least one subsequent instruction that is not yet ready for execution. However, Blomgren has taught such a concept. See column 5, line 67, to column 6, line 2, and column 7, lines 3-7. Note that up to 3 instructions may be dispatched (issued) at the same time. However, if a stall occurs in one of the three pipelines (each one for executing one of the three instructions), then all three pipelines are stalled and zero instructions are issued. So, for example, if three instructions are set to be dispatched, and of the three instructions, two are considered ready (no dependencies on previous instructions), but one is not ready because of a dependency issue, then all three pipelines will be stalled, causing the issue of the two ready instructions to be delayed until the dependency problem is taken care of. As disclosed in column 5, line 67, to column 6, line 2, such a stalling mechanism simplifies the system. As a result, it would have been obvious to one of ordinary skill in the art at the time of

Art Unit: 2183

the invention to modify Col, etc. to delay the issue of a first instruction if it is determined that said first microinstruction is required to be issued in parallel with at least one subsequent instruction that is not yet ready for execution.

Conclusion

37. Applicant's amendment necessitated the new ground(s) of rejection presented in this Office action. Accordingly, **THIS ACTION IS MADE FINAL**. See MPEP § 706.07(a). Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).

A shortened statutory period for reply to this final action is set to expire **THREE MONTHS** from the mailing date of this action. In the event a first reply is filed within **TWO MONTHS** of the mailing date of this final action and the advisory action is not mailed until after the end of the **THREE-MONTH** shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of the advisory action. In no event, however, will the statutory period for reply expire later than **SIX MONTHS** from the date of this final action.

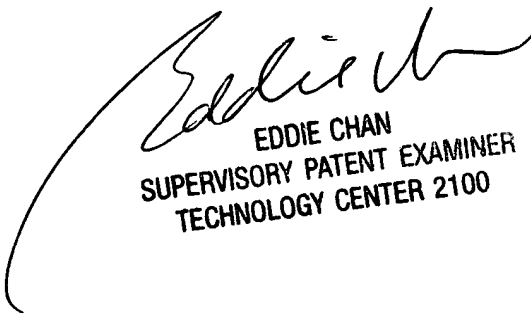
Any inquiry concerning this communication or earlier communications from the examiner should be directed to David J. Huisman whose telephone number is (703) 305-7811. The examiner can normally be reached on Monday-Friday (8:00-4:30).

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Eddie Chan can be reached on (703) 305-9712. The fax phone number for the organization where this application or proceeding is assigned is 703-872-9306.

Art Unit: 2183

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

DJH
David J. Huisman
April 14, 2004



EDDIE CHAN
SUPERVISORY PATENT EXAMINER
TECHNOLOGY CENTER 2100